

# The complexity of computation in bit streams

Raphaël Clifford    Markus Jalsenius    Benjamin Sach

Department of Computer Science  
University of Bristol  
Bristol, UK

## Abstract

We revisit the complexity of online computation in the cell probe model. We consider a class of problems where we are first given a fixed pattern or vector  $F$  of  $n$  symbols and then one symbol arrives at a time in a stream. After each symbol has arrived we must output some function of  $F$  and the  $n$ -length suffix of the arriving stream. Cell probe bounds of  $\Omega(\delta \lg n / w)$  have previously been shown for both convolution and Hamming distance in this setting, where  $\delta$  is the size of a symbol in bits and  $w \in \Omega(\lg n)$  is the cell size in bits. However, when  $\delta$  is a constant, as it is in many natural situations, these previous results no longer give us non-trivial bounds.

We introduce a new *lop-sided information transfer* proof technique which enables us to prove meaningful lower bounds even for constant size input alphabets. We use our new framework to prove an amortised cell probe lower bound of  $\Omega(\lg^2 n / (w \cdot \lg \lg n))$  time per arriving bit for an online version of a well studied problem known as pattern matching with address errors. This is the first non-trivial cell probe lower bound for any online problem on bit streams that still holds when the cell sizes are large. We also show the same bound for online convolution conditioned on a new combinatorial conjecture related to Toeplitz matrices.

## 1 Introduction

We revisit the complexity of online computation in the cell probe model. In recent years there has been considerable progress towards the challenging goal of establishing lower bounds for both static and dynamic data structure problems. A third class of data structure problems, one which falls somewhere between these two classic settings, is online computation in a streaming setting. Here one symbol arrives at a time and a new output must be given after each symbol arrives and before the next symbol is processed. The key conceptual difference to a standard dynamic data structure problem is that although each arriving symbol can be regarded as a new update operation, there is only one type of query which is to output the latest value of some function of the stream.

Online pattern matching is particularly suited to study in this setting and cell probe lower bounds have previously been shown for different measures of distance including Hamming distance, inner product/convolution and edit distance [5, 6, 4]. All these previous cell probe lower bounds have relied on only one proof technique, the so-called *information transfer technique* of Pătraşcu and Demaine [15]. In loose terms the basic idea is as follows. First one defines a random input distribution over updates. Here we regard an arriving symbol as an update and after each update we perform one query which simply returns the latest distance between a predefined pattern and the updated suffix of the stream. Then one has to argue that knowledge of the answers to  $\ell$  consecutive queries is sufficient to infer at least a constant fraction of the information encoded by  $\ell$  consecutive updates that occurred in the past. If one can show this is true for all power of two lengths  $\ell$  and ensure there is no double counting then the resulting lower bound follows by summing over all these power of two lengths.

For the most natural cell size  $w \in \Omega(\lg n)$ , a cell probe lower bound of  $\Omega(\delta \lg n/w)$  for both Hamming distance and convolution using this method was shown, where  $\delta$  is the size of an input symbol,  $w$  is the cell size in bits and  $n$  is the length of the fixed pattern [5, 6]. When  $\delta = w$ , there is also a matching upper bound in the cell probe model and so no further progress is possible. However, when the symbol size  $\delta$  is in fact constant, the best lower bound that is derivable reduces trivially to be constant. This is a particularly unfortunate situation as arguably the *most* natural setting of parameters is when the input alphabet is of constant size but the cell size is not.

To make matters worse, this limitation is neither specific to pattern matching problems nor even to online problems in general. As it is a fundamental feature of the information transfer technique the requirement to have large input alphabets also applies to a wide of dynamic data structure problems for which the information transfer technique has been up to this point the lower bound method of choice. As a result we see the challenge of providing a new proof technique which can meaningfully handle constant sized alphabets as fundamental to the aim of advancing our knowledge of the true complexity of both online and dynamic problems.

We introduce a new lop-sided version of the information transfer technique that enables us to give meaningful lower bounds for precisely this setting, that is when  $\delta \in O(1)$  and  $w \in \Omega(\lg n)$ . Our proof technique will rely on being able to show for specific problems that we need only  $\ell$  query answers to give at least a constant fraction of the information encoded in  $\ell \lg \ell$  updates.

We demonstrate our new framework by first by applying it to the well studied online convolution problem. For this problem we give a conditional cell probe lower bound which depends on a new combinatorial conjecture involving Toeplitz matrices. We then show that it is possible to derive an identical but this time unconditional lower bound by applying the lop-sided information transfer technique to a problem called online pattern matching with address errors [2]. This measure of distance arises in pattern matching problems where errors occur not in the content of the data but in the addresses for where the data is stored.

## Previous cell probe lower bounds

Our bounds hold in a particularly strong computational model, the *cell-probe model*, introduced originally by Minsky and Papert [12] in a different context and then subsequently by Fredman [8] and Yao [17]. In this model, there is a separation between the computing unit and the memory, which is external and consists of a set of cells of  $w$  bits each. The computing unit cannot remember any information between operations. Computation is free and the cost is measured only in the number of cell reads or writes (cell-probes). This general view makes the model very strong, subsuming for instance the popular word-RAM model.

The first techniques known for establishing dynamic data structure lower bounds had historically been based on the chronogram technique of Fredman and Saks [9] which can at best give bounds of  $\Omega(\lg n / \lg \lg n)$ . In 2004, Pătraşcu and Demaine gave us the first  $\Omega(\lg n)$  lower bounds for dynamic data structure problems [15]. Their technique is based on information theoretic arguments which also form the basis for the work we present in this paper. Pătraşcu and Demaine also presented ideas which allowed them to express more refined lower bounds such as trade-offs between updates and queries of dynamic data structures. For a list of data structure problems and their lower bounds using these and related techniques, see for example [13]. More recently, a further breakthrough was made by Larsen who showed lower bounds of roughly  $\Omega((\lg n / \lg \lg n)^2)$  time per operation for dynamic weighted range counting problem and polynomial evaluation [10, 11]. These lower bounds remain the state of the art for any dynamic structure problem to this day.

### 1.1 Our Results

**The lop-sided information transfer technique** In the standard formulation of the information transfer technique of Demaine and Pătraşcu [14], adjacent time intervals are considered and the *information* that is transferred from the operations in one interval of power of two length  $\ell$  to the next interval of the same length is studied. It is sufficient to prove that a lower bound can be given for the information transfer which applies to all consecutive intervals of power of two length. Conceptually a balanced tree on  $n$  leaves over the time axis is constructed which is known as the information transfer tree. An internal node  $v$  is associated with the times  $t_0$ ,  $t_1$  and  $t_2$  such that the two intervals  $[t_0, t_1]$  and  $[t_1 + 1, t_2]$  span the left subtree and the right subtree of  $v$ , respectively. By summing over the information transfer at each node  $v$  the final lower bound is derived.

In order to show a cell probe lower bound for bit streams, we will need to give lower bounds for the information transferred from intervals of time  $[t_0, t_1]$  to later intervals  $[t_2, t_3]$  which are shorter than the first intervals. In particular we will want to argue about intervals of length  $\ell \lg \ell$  and  $\ell$  respectively. However, making the interval lengths lop-sided requires us to abandon the information transfer tree and indeed almost all of the previous proof technique. We will instead place gaps in time between the end of one interval and the start

of another and argue carefully both that not too much of the information can be lost in these gaps and that we can still sum the information transfer over a sufficient number of distinct interval lengths without too much double counting.

Our hope is that this new technique will lead to a new class of cell probe lower bounds which could not be proved with existing methods.

**Online convolution** In the online convolution problem we are given a fixed vector  $F \in \{0, 1\}^n$  and a stream that arrives one bit at a time. After each bit arrives we must output the inner product between  $F$  and a vector formed from the most recent  $n$ -length suffix of the stream. At the heart of our lower bound we need to show that we can recover  $\Omega(\ell \lg \ell)$  bits of information of a contiguous subarray of the stream of length  $\ell \lg \ell$ . However, because the inputs are binary and to avoid a trivial lower bound, we must do so using only  $\ell$  outputs. As we need to achieve this for a large number distinct values of  $\ell$  simultaneously, the fixed vector  $F$  we design for the online convolution problem has a carefully designed recursive structure. However even having overcome this hurdle, there is still the challenge of showing that for each separate value  $\ell$ , the information from successive outputs does not have too large an overlap. A complete resolution to this question seems non-trivial but in Section 4 we present a new combinatorial conjecture related to Toeplitz matrices, which if true, provides the desired lower bound. As a result we get the following lower bound for online convolution in a bit stream:

**Theorem 1 (Online convolution).** *Assuming Conjecture 1, in the cell-probe model with  $w \in \Omega(\lg n)$  bits per cell, for any randomised algorithm solving the online convolution problem on binary inputs there exist instances such that the expected amortised time per arriving value is*

$$\Omega\left(\frac{\lg^2 n}{w \cdot \lg \lg n}\right).$$

**Online pattern matching with address errors ( $L_2$ -rearrangement distance)** As our second example we give an explicit distance function for which we can now obtain the first unconditional online lower bound for symbol size  $\delta = 1$ . We consider a problem setting known as pattern matching with address errors and in particular the  $L_2$ -rearrangement distance as defined in [2]. Consider two strings  $S_1$  and  $S_2$  both of length  $n$  where  $S_2$  is a permutation of  $S_1$ . Now consider the set of permutations  $\Pi$  so that for all  $\pi \in \Pi$ ,  $S_1[\pi(0), \dots, \pi(n-1)] = S_2$ . The  $L_2$ -rearrangement distance is defined to be  $\min_{\pi \in \Pi} \sum_{j=0}^{n-1} |j - \pi(j)|^2$ . If  $\Pi$  is empty, that is  $S_2$  is in fact not a permutation of  $S_1$ , then the  $L_2$ -rearrangement distance is defined to be  $\infty$ . When considered as an offline problem, for a text of length  $2n$  and a pattern of length  $n$ , Amir et al. showed that the  $L_2$ -rearrangement distance between the pattern and every  $n$ -length substring of the text could be computed in  $O(n \lg n)$  time [2]. In the online  $L_2$ -rearrangement problem we are given a fixed pattern  $F \in \{0, 1\}^n$  and the stream arrives one bit at a time. After each

bit arrives we must output the  $L_2$ -rearrangement distance between  $F$  and the most recent  $n$ -length suffix of the stream.

As before we need to recover  $\Omega(\ell \lg \ell)$  bits of information of a contiguous sub-array of the stream of length  $\ell \lg \ell$ . Our technique allows us to recover  $\Omega(\lg n)$  distinct bits of stream from each output. This is achieved by constructing  $F$  and carefully choosing a highly structured random input distribution for the incoming stream in such a way that the contributions to the output from different regions of the stream have different magnitudes. We can then use the result to extract distinct information about the stream from different parts of each output.

Using this approach we get the following cell probe lower bound:

**Theorem 2 (Online  $L_2$ -rearrangement).** *In the cell-probe model with  $w \in \Omega(\lg n)$  bits per cell, for any randomised algorithm solving the online  $L_2$ -rearrangement distance problem on binary inputs there exist instances such that the expected amortised time per arriving value is*

$$\Omega\left(\frac{\lg^2 n}{w \cdot \lg \lg n}\right).$$

## 2 Lop-sided information transfer

We will define the concept of information transfer, a particular set of cells probed by the algorithm, and explain how a bound on the size of the information transfer can be used when proving the overall lower bounds of Theorems 1 and 2. All logarithms are in base two.

### 2.1 Notation for the online problems

We define some notation for our online problems. There is a *fixed array*  $F \in \{0, 1\}^n$  of length  $n$  and an array  $S \in \{0, 1\}^n$  of length  $n$ , referred to as the *stream*. We maintain  $S$  subject to an update operation  $\text{UPDATE}(x)$  which takes a value  $x \in \{0, 1\}$ , modifies  $S$  by appending  $x$  to the right of the rightmost component  $S[n - 1]$  and removing the leftmost component  $S[0]$ , and then outputs the inner product of  $F$  and  $S$ , that is  $\sum_{i \in [n]} (F[i] \cdot S[i])$ , or alternatively the  $L_2$ -rearrangement distance, depending on which problem we are currently considering.

We let  $U \in \{0, 1\}^n$  denote the *update array* which describes a sequence of  $n$   $\text{UPDATE}$  operations. That is, for each  $t \in [n]$ , the operation  $\text{UPDATE}(U[t])$  is performed. We will usually refer to  $t$  as the *arrival* of the value  $U[t]$ . Observe that just after arrival  $t$ , the values  $U[t + 1, n - 1]$  are still not known to the algorithm. Finally, we let the  $n$ -length array  $A$  denote the outputs such that for  $t \in [n]$ ,  $A[t]$  is the output of  $\text{UPDATE}(U[t])$ .

## 2.2 Hard distributions

Our lower bound holds for any randomised algorithm on its worst case input. This will be achieved by applying *Yao's minimax principle* [16]. We develop a lower bound that holds for any deterministic algorithm on some random input. The basic approach is as follows: we devise a fixed array  $F$  and describe a probability distribution for  $U$ , the  $n$  new values arriving in the stream  $S$ . We then obtain a lower bound on the expected running time over the arrivals in  $U$  that holds for any deterministic algorithm. Due to the minimax principle, the same lower bound must then also hold for any randomised algorithm on its own worst case input. The amortised bound per arriving value is obtained by dividing by  $n$ .

From this point onwards we consider an arbitrary deterministic algorithm running with some fixed array  $F$  on a random input of  $n$  values. The algorithm may depend on  $F$ . We refer to the choice of  $F$  and distribution for  $U$  as a *hard distribution* since it is used to show a lower bound.

## 2.3 Two intervals and a gap

In order to define the concept of information transfer from one interval of arriving values in the stream to another interval of arriving values, we first define the set  $L$  which contains the interval lengths that we will consider. The purpose of the next few definitions will be clear once we define information transfer in the next section. We let

$$L = \left\{ n^{1/4} \cdot (\lg n)^{2i} \mid i \in \left\{ 0, 1, 2, \dots, \frac{\lg n}{4 \lg \lg n} \right\} \right\}.$$

To avoid cluttering the presentation with floors and ceilings, we assume throughout that the value of  $n$  is such that any division or power nicely yields an integer. Whenever it is impossible to obtain an integer we assume that suitable floors or ceilings are used. In particular,  $L$  contains only integers.

For  $\ell \in L$  and  $t \in [n/2]$  we define the following four values:

$$\begin{aligned} t_0 &= t, \\ t_1 &= t_0 + \ell \lg \ell - 1, \\ t_2 &= t_1 + \frac{4\ell}{\lg n} + 1, \\ t_3 &= t_2 + \ell - 1. \end{aligned}$$

The values  $t_0$ ,  $t_1$ ,  $t_2$  and  $t_3$  are indeed functions of  $\ell$  and  $t$  but for brevity we will often write just  $t_0$  instead of  $t_0(\ell, t)$ , and so on, whenever the parameters  $\ell$  and  $t$  are obvious from context. The four values define the three intervals  $[t_0, t_1]$ ,  $[t_1 + 1, t_2 - 1]$  and  $[t_2, t_3]$ , referred to as the *first interval*, the *gap* and the *second interval*, respectively. Whenever

there is a risk of ambiguity of which parameters the intervals are based on, we may specify what  $\ell$  and  $t$  they are associated with. Before we explain the purpose of the intervals we will highlight some of their properties.

First observe that the intervals are disjoint and the gap has length  $4\ell/\lg n$ . The first interval has length  $\ell \lg \ell$  and starts at  $t$ , where  $t$  is always in the first half of the interval  $[0, n-1]$ . The second interval has length  $\ell$ , hence is a log-factor shorter than the first interval. All intervals are contained in  $[0, n-1]$ . To see this we need to verify that  $t_3 \leq n-1$  for any choice of  $\ell \in L$  and  $t \in [n/2]$ . The largest value in  $L$  is

$$n^{1/4} \cdot (\lg n)^{\frac{2 \lg n}{4 \lg \lg n}} = n^{3/4}.$$

Thus, the largest possible value of  $t_3$ , obtained with  $\ell = n^{3/4}$  and  $t = n/2 - 1$ , is

$$\begin{aligned} \left(\frac{n}{2} - 1\right) + \left(n^{3/4} \cdot \lg n^{3/4} - 1\right) + \left(\frac{4n^{3/4}}{\lg n} + 1\right) + \left(n^{3/4} - 1\right) \\ \leq \frac{n}{2} + 1 + n^{3/4}(\lg n^{3/4} + 5) < n, \end{aligned}$$

whenever  $n$  is sufficiently large.

Lastly, suppose that  $\ell' \in L$  is one size larger than  $\ell \in L$ , that is  $\ell' = \ell \cdot (\lg n)^2$ . For  $\ell'$  the length of the gap is  $4\ell'/\lg n$ , which so big that it spans the length of both intervals plus the gap associated with  $\ell$ . To see this, observe that for sufficiently large  $n$ ,

$$\begin{aligned} t_3(\ell, t) - t_0(\ell, t) + 1 &= \ell \lg \ell - 1 + \frac{4\ell}{\lg n} + 1 + \ell - 1 + 1 \\ &\leq 2\ell \lg n + \frac{4\ell}{\lg n} \leq 2\ell \lg n + \frac{(\lg n)^2 \ell}{\lg n} = 3\ell \lg n \leq \frac{4\ell'}{\lg n}. \end{aligned}$$

## 2.4 Information transfer over gaps

Towards the definition of information transfer, we define, for  $\ell \in L$  and  $t \in [n/2]$ , the subarray  $U_{\ell,t} = U[t_0, \dots, t_1]$  to represent the  $\ell \lg \ell$  values arriving in the stream during the first interval. We define the subarray  $A_{\ell,t} = A[t_2, \dots, t_3]$  to represent the  $\ell$  outputs during the second interval. Lastly we define  $\tilde{U}_{\ell,t}$  to be the concatenation of  $U[0, (t_0 - 1)]$  and  $U[(t_1 + 1), (n - 1)]$ . That is,  $\tilde{U}_{\ell,t}$  contains all values of  $U$  except for those in  $U_{\ell,t}$ .

For  $\ell \in L$  and  $t \in [n/2]$  we first define the *information transfer to the gap*, denoted  $\mathcal{G}_{\ell,t}$ , to be the set of memory cells  $c$  such that  $c$  is probed during the first interval  $[t_0, t_1]$  of arriving values and also probed during the arrivals of the values  $U[t_1 + 1, t_2 - 1]$  in the gap. Similarly we define the information transfer to the second interval, or simply *the information transfer*, denoted  $\mathcal{I}_{\ell,t}$ , to be the set of memory cells  $c$  such that  $c$  is probed during the first interval  $[t_0, t_1]$  of arriving symbols and also probed during the arrivals of symbols in the second interval  $[t_2, t_3]$  but *not* in the gap. That is, any cell  $c \in \mathcal{G}_{\ell,t}$  cannot also be contained in the information transfer  $\mathcal{I}_{\ell,t}$ .

The cells in the information transfer  $\mathcal{I}_{\ell,t}$  may contain information about the values in  $U_{\ell,t}$  that the algorithm uses in order to correctly produce the outputs  $A_{\ell,t}$ . However, since cells that are probed in the gap are not included in the information transfer, the information transfer might not contain all the information about the values in  $U_{\ell,t}$  that the algorithm uses while outputting  $A_{\ell,t}$ . In all previous work on lower bounds where the information transfer technique is used, the two intervals had no gap between them, hence the information transfer  $\mathcal{I}_{\ell,t}$  contained *all* the information about the updates in  $U_{\ell,t}$  necessary for computing  $A_{\ell,t}$ . Further, in all previous work, the two intervals always had the same length. We will see that the introduction of a gap and skewed interval lengths enable us to provide non-trivial lower bounds for small inputs with large cell sizes. We will see that the gap is small enough that a sufficiently large portion of the information about  $U_{\ell,t}$  has to be fetched from cells in the information transfer  $\mathcal{I}_{\ell,t}$ .

Since cells in the information transfer are by definition probed at some point by the algorithm, we can use  $\mathcal{I}_{\ell,t}$  to measure, or at least lower bound, the number of cell probes. As a shorthand we let  $I_{\ell,t} = |\mathcal{I}_{\ell,t}|$  denote the size of the information transfer  $\mathcal{I}_{\ell,t}$ . Similarly we let  $G_{\ell,t} = |\mathcal{G}_{\ell,t}|$  denote the size of the information transfer to the gap. By adding up the sizes  $I_{\ell,t}$  of the information transfers over all  $\ell \in L$  and certain values of  $t \in [n/2]$ , we get a lower bound on the total number of cells probed by the algorithm during the  $n$  arriving values in  $U$ . The choice of the values  $t$  is crucial as we do not want to over-count the number of cell probes. In the next two lemmas we will deal with the potential danger of over-counting.

For a cell  $c \in \mathcal{I}_{\ell,t}$ , we write *the probe of  $c$  with respect to  $\mathcal{I}_{\ell,t}$*  to refer to the first probe of  $c$  during the arrivals in the second interval. These are the probes of the cells in the information transfer that we count.

**Lemma 1.** *For any  $\ell \in L$  and  $t, t' \in [n/2]$  such that  $|t - t'| \geq \ell$ , if a cell  $c$  is in both  $\mathcal{I}_{\ell,t}$  and  $\mathcal{I}_{\ell,t'}$  then the probe of  $c$  with respect to  $\mathcal{I}_{\ell,t}$  and the probe of  $c$  with respect  $\mathcal{I}_{\ell,t'}$  are distinct.*

*Proof.* Since  $t$  and  $t'$  are at least  $\ell$  apart, the second intervals associated with  $t$  and  $t'$ , respectively, must be disjoint. Hence the probe of  $c$  with respect to  $\mathcal{I}_{\ell,t}$  and the probe of  $c$  with respect  $\mathcal{I}_{\ell,t'}$  must be distinct.  $\square$

From the previous lemma we know that there is no risk of over-counting cell probes of a cell over information transfers  $\mathcal{I}_{\ell,t}$  under a fixed value of  $\ell \in L$ , as long as no two values of  $t$  are closer than  $\ell$ . In the next lemma we consider information transfers under different values of  $\ell \in L$ .

**Lemma 2.** *For any  $\ell, \ell' \in L$  such that  $\ell \neq \ell'$ , and any  $t, t' \in [n/2]$ , if a cell  $c$  is in both  $\mathcal{I}_{\ell,t}$  and  $\mathcal{I}_{\ell',t'}$  then the probe of  $c$  with respect to  $\mathcal{I}_{\ell,t}$  and the probe of  $c$  with respect  $\mathcal{I}_{\ell',t'}$  must be distinct.*



*Proof.* Let  $p$  be the probe of  $c$  with respect to  $\mathcal{I}_{\ell,t}$ , and let  $p'$  be the probe of  $c$  with respect to  $\mathcal{I}_{\ell',t'}$ . We will show that  $p \neq p'$ . Suppose without loss of generality that  $\ell < \ell'$ . From the properties of the intervals that were given in the previous section we know that the length of the gap associated with  $\ell'$  is larger than the sum of lengths of the first interval, the gap and the second interval associated with  $\ell$ .

Suppose for contradiction that  $p = p'$ . By definition of  $\mathcal{I}_{\ell,t}$ , the cell  $c$  is probed also in the first interval associated with  $\ell$ . Let  $p_{\text{first}}$  denote any such cell probe. Because the gap associated with  $\ell'$  is so large,  $p_{\text{first}}$  must take place either in the second interval or the gap associated with  $\ell'$ . If  $p_{\text{first}}$  is in the gap, then  $c$  cannot be in  $\mathcal{I}_{\ell',t'}$ . If  $p_{\text{first}}$  is in the second interval then  $p'$  cannot equal  $p$ .  $\square$

In order to lower bound the total number of cell probes performed by the algorithm over the  $n$  arrivals in  $U$  we will define, for each  $\ell \in L$ , a set  $T_\ell \subseteq [n/2]$  of arrivals, such that for any distinct  $t, t' \in T_\ell$ ,  $|t - t'| \geq \ell$ . It then follows from Lemmas 1 and 2 that

$$\sum_{\ell \in L} \sum_{t \in T_\ell} I_{\ell,t}$$

is a lower bound on the number of cell probes. Our goal is to lower bound the expected value of this double-sum. The exact definition of  $T_\ell$  will be given in Section 3.3 once we have introduced relevant notation.

### 3 Proving the lower bound

In this section we give the overall proof for the lower bounds of Theorems 1 and 2. Let  $\ell \in L$  and let  $t \in [n/2]$ . Suppose that  $\tilde{U}_{\ell,t}$  is fixed but the values in  $U_{\ell,t}$  are drawn at random in accordance with the distribution for  $U$ , conditioned on the fixed value of  $\tilde{U}_{\ell,t}$ . This induces a distribution for the outputs  $A_{\ell,t}$ . We want to show that if the entropy of  $A_{\ell,t}$  is large, conditioned on the fixed  $\tilde{U}_{\ell,t}$ , then the information transfer  $\mathcal{I}_{\ell,t}$  is large, since only the variation in the inputs  $U_{\ell,t}$  can alter the outputs  $A_{\ell,t}$ . We will soon make this claim more precise.

#### 3.1 Upper bound on the entropy

We write  $H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t})$  to denote the entropy of  $A_{\ell,t}$  conditioned on fixed  $\tilde{U}_{\ell,t}$ . Towards showing that high conditional entropy  $H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t})$  implies large information transfer we use the information transfer  $\mathcal{I}_{\ell,t}$  and the information transfer to the gap,  $\mathcal{G}_{\ell,t}$ , to describe an encoding of the outputs  $A_{\ell,t}$ . The following lemma gives a direct relationship between  $I_{\ell,t} + G_{\ell,t}$  and the entropy. A marginally simpler version of the lemma, stated with different notation, was first given in [15] under the absence of gaps.

**Lemma 3.** *Under the assumption that the address of any cell can be specified in  $w$  bits, for any  $\ell \in L$  and  $t \in [n/2]$ , the entropy*

$$H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t}) \leq 2w + 2w \cdot \mathbb{E}[I_{\ell,t} + G_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t}].$$

*Proof.* The expected length of any encoding of  $A_{\ell,t}$  under fixed  $\tilde{U}_{\ell,t}$  is an upper bound on the conditional entropy of  $A_{\ell,t}$ . We use the information transfer  $\mathcal{I}_{\ell,t}$  and the information transfer to the gap,  $\mathcal{G}_{\ell,t}$ , to define an encoding of  $A_{\ell,t}$  in the following way. For every cell  $c \in \mathcal{I}_{\ell,t} \cup \mathcal{G}_{\ell,t}$  we store the address of  $c$ , which takes at most  $w$  bits under the assumption that a cell can hold the address of any cell in memory. We also store the contents of  $c$  that it holds at the very end of the first interval, just before the beginning of the gap. The contents of  $c$  is specified with  $w$  bits. In total this requires  $2w \cdot (I_{\ell,t} + G_{\ell,t})$  bits.

We will use the algorithm, which is fixed, and the fixed values  $\tilde{u}_{\ell,t}$  of  $\tilde{U}_{\ell,t}$  as part of the decoder to obtain  $A_{\ell,t}$  from the encoding. Since the encoding is of variable length we also store the size  $I_{\ell,t}$  of the information transfer and the size  $G_{\ell,t}$  of the information transfer to the gap. This requires at most  $2w$  additional bits.

In order to prove that the described encoding of  $A_{\ell,t}$  is valid we now describe how to decode it. First we simulate the algorithm on the fixed input  $\tilde{U}_{\ell,t}$  from the first arrival  $U[0]$  until just before the first interval when the first value in  $U_{\ell,t}$  arrives. We then skip over all inputs in  $U_{\ell,t}$  and resume simulating the algorithm from the beginning of the gap, that is when the value  $U[t_1 + 1]$  arrives. We simulate the algorithm over the arrivals in the gap and the second interval until all values in  $A_{\ell,t}$  have been outputted. For every cell being read, we check if it is contained in either the information transfer  $\mathcal{I}_{\ell,t}$  or the information transfer to the gap  $\mathcal{G}_{\ell,t}$  by looking up its address in the encoding. If the address is found then the contents of the cell is fetched from the encoding. If not, its contents is available from simulating the algorithm on the fixed inputs  $\tilde{U}_{\ell,t}$ .  $\square$

### 3.2 Lower bounds on entropy

Lemma 3 above provides a direct way to obtain a lower bound on the expected value of  $I_{\ell,t} + G_{\ell,t}$  if given a lower bound on the conditional entropy  $H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t})$ . In the next two lemmas we provide such entropy lower bounds for  $L_2$ -rearrangement distance and convolution.

**Lemma 4.** *Assuming Conjecture 1, for the convolution problem there exists a real constant  $\kappa > 0$  and, for any  $n$ , a fixed array  $F \in \{0, 1\}^n$  such that for all  $\ell \in L$  and  $t \in [n/2]$ , when  $U$  is chosen uniformly at random from  $\{0, 1\}^n$  then*

$$H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t}) \geq \kappa \cdot \ell \cdot \lg n,$$

*for any fixed  $\tilde{u}_{\ell,t}$ .*

**Lemma 5.** *For the  $L_2$ -rearrangement distance problem there exists a real constant  $\kappa > 0$  and, for any  $n$ , a fixed array  $F \in \{0, 1\}^n$  such that for all  $\ell \in L$  and all  $t \in [n/2]$  such that  $t \bmod 4 = 0$ , when  $U$  is chosen uniformly at random from  $\{0101, 1010\}^{\frac{n}{4}}$  then*

$$H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t}) \geq \kappa \cdot \ell \cdot \lg n,$$

for any fixed  $\tilde{u}_{\ell,t}$ .

The proof of Lemma 4 is deferred to Section 4 and hinges on a conjecture relating to Toeplitz matrices. The proof of Lemma 5 is deferred to Section 5.

Before we proceed with the lower bound on the information transfer we make a short remark on the bounds that these lemmas give. Observe that the maximum conditional entropy of  $A_{\ell,t}$  is bounded by the entropy of  $U_{\ell,t}$ , which is  $O(\ell \lg \ell)$  since the length of the first interval is  $\ell \lg \ell$ . Recall also that the values in  $L$  range from  $n^{1/4}$  to  $n^{3/4}$ . Thus, for a constant  $\kappa$ , both entropy lower bounds are tight up to a multiplicative constant factor.

### 3.3 A lower bound on the information transfer and quick gaps

In this section we prove our main lower bound results. In order to fix ideas we do this for the  $L_2$ -rearrangement problem, given in Theorem 2. As a result we assume that  $\kappa$  is the constant and  $F$  is the fixed array of Lemma 4, and that  $U$  is chosen uniformly at random from  $\{0, 1\}^n$ . Theorem 1 which gives our main lower bound result for online convolution follows via exactly the same argument but with Lemma 4 replaced by Lemma 5.

By combining the upper and lower bounds on the conditional entropy from Lemmas 3 and 5 we have that there is a hard distribution and a real constant  $\kappa > 0$  for the convolution problem such that

$$\mathbb{E}[I_{\ell,t} + G_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t}] \geq \frac{\kappa \cdot \ell \cdot \lg n}{2w} - 1$$

for any  $\tilde{u}_{\ell,t}$ . We may remove the conditioning by taking expectation over  $\tilde{U}_{\ell,t}$  under random  $U$ . Thus,

$$\mathbb{E}[I_{\ell,t} + G_{\ell,t}] \geq \frac{\kappa \cdot \ell \cdot \lg n}{2w} - 1,$$

or equivalently

$$\mathbb{E}[I_{\ell,t}] \geq \frac{\kappa \cdot \ell \cdot \lg n}{2w} - 1 - \mathbb{E}[G_{\ell,t}]. \quad (1)$$

Recall that our goal is to lower bound

$$\mathbb{E} \left[ \sum_{\ell \in L} \sum_{t \in T_\ell} I_{\ell,t} \right] = \sum_{\ell \in L} \sum_{t \in T_\ell} \mathbb{E}[I_{\ell,t}],$$

where  $T_\ell$  contains suitable values of  $t$ . Using Inequality (1) above would immediately provide such a lower bound, however, there is an imminent risk that the negative terms of

$\mathbb{E}[G_{\ell,t}]$  could devalue such a bound into something trivially small. Now, for this to happen, the algorithm must perform sufficiently many cell probes in the gap. Since the length of the gap is considerably shorter than the second interval, a cap on the worst-case number of cell probes per arriving value would certainly ensure that  $\mathbb{E}[G_{\ell,t}]$  stays small, but as we want an amortised lower bound we need something more refined. The answer lies in how we define  $T_\ell$ . We discuss this next.

For  $\ell \in L$  and  $f \in [\ell]$  we first define the set  $T_{\ell,f} \subseteq [n]$  of arrivals to be

$$T_{\ell,f} = \left\{ f + i\ell \mid i \in \{0, 1, 2, \dots\} \text{ and } f + i\ell \leq \frac{n}{2} \right\}.$$

The values in  $T_{\ell,f}$  are evenly spread out with distance  $\ell$ , starting at  $f$ . We may think of  $f$  as the offset of the sequence of values in  $T_{\ell,f}$ . The largest value in the set is no more than  $n/2$ . We will define the set  $T_\ell$  to equal a subset of one of the sets  $T_{\ell,f}$  for some  $f$ . More precisely, we will show that there must exist an offset  $f$  such that at least half of the values  $t \in T_{\ell,f}$  have the property that the time spent in the gap associated with  $\ell$  and  $t$  is small enough to ensure that the information transfer to the gap is small. We begin with some definitions.

**Definition 1 (Quick gaps and sets).** For any  $\ell \in L$  and  $t \in [n/2]$  we say that the gap associated with  $\ell$  and  $t$  is *quick* if the expected number of cell probes during the arrivals in the gap is no more than

$$\frac{\kappa \cdot \ell \cdot \lg n}{4w},$$

where  $\kappa$  is the constant from Lemma 4. Further, for any  $f \in [\ell]$  we say that the set  $T_{\ell,f}$  is *quick* if, for at least half of all  $t \in T_{\ell,f}$ , the gap associated with  $\ell$  and  $t$  is quick.

The next lemma says that for sufficiently fast algorithms there is always an offset  $f$  such that  $T_{\ell,f}$  is quick.

**Lemma 6.** *Suppose that the expected total number of cell probes over the  $n$  arrivals in  $U$  is less than*

$$\frac{\kappa \cdot n \cdot \lg^2 n}{32w}.$$

*Then, for any  $\ell \in L$ , there is an  $f \in [\ell]$  such that  $T_{\ell,f}$  is quick.*

*Proof.* In accordance with the lemma, suppose that the expected total number of cell probes over the  $n$  arrivals in  $U$  is less than  $\kappa n(\lg^2 n)/(32w)$ . For contradiction, suppose that there is no  $f \in [\ell]$  such that  $T_{\ell,f}$  is quick. We will show that the expected number of cell probes over the  $n$  arrivals must then be at least  $\kappa n(\lg^2 n)/(32w)$ .

For any  $f \in [\ell]$ , let  $R_f \subseteq [n]$  be the union of all arrivals that belong to a gap associated with  $\ell$  and any  $t \in T_{\ell,f}$ . Let  $P_f$  be the number of cell probes performed by the algorithm

over the arrivals in  $R_f$ . Thus, for any set  $T_{\ell,f}$  that is *not quick* we have by linearity of expectation

$$\mathbb{E}[P_f] \geq \frac{|T_{\ell,f}|}{2} \cdot \frac{\kappa \cdot \ell \cdot \lg n}{4w} = \frac{n/2}{2\ell} \cdot \frac{\kappa \cdot \ell \cdot \lg n}{4w} = \frac{\kappa \cdot n \cdot \lg n}{8w}.$$

Let the set of offsets

$$\mathcal{F} = \left\{ i \cdot \frac{4\ell}{\lg n} \mid i \in \left[ \frac{\lg n}{4} \right] \right\} \subseteq [\ell].$$

The values in  $\mathcal{F}$  are spread out with distance  $4\ell/\lg n$ , which equals the gap length. Thus, for any two distinct  $f, f' \in \mathcal{F}$ , the sets  $R_f$  and  $R_{f'}$  are disjoint. We therefore have that the total running time over all  $n$  arrivals in  $U$  must be lower bounded by  $\sum_{f \in \mathcal{F}} P_f$ . Under the assumption that no  $T_{\ell,f}$  is quick, we have that the expected total running time is at least

$$\mathbb{E} \left[ \sum_{f \in \mathcal{F}} P_f \right] = \sum_{f \in \mathcal{F}} \mathbb{E}[P_f] \geq |\mathcal{F}| \cdot \frac{\kappa \cdot n \cdot \lg n}{8w} = \frac{\lg n}{4} \cdot \frac{\kappa \cdot n \cdot \lg n}{8w} = \frac{\kappa \cdot n \cdot \lg^2 n}{32w},$$

which is the contradiction we wanted. Thus, under the assumption that the running time over the  $n$  arrivals in  $U$  is less than  $\kappa n (\lg^2 n) / (32w)$  there must be an  $f \in [\ell]$  such that  $T_{\ell,f}$  is quick.  $\square$

We now proceed under the assumption that the expected running time over the  $n$  arrivals in  $U$  is less than  $\kappa n (\lg^2 n) / (32w)$ . If this is not the case then we have already established the lower bound of Theorem 1.

Let  $f$  be a value in  $[\ell]$  such that  $T_{\ell,f}$  is a quick set. Such an  $f$  exists due to Lemma 6. We now let  $T_\ell \subseteq T_{\ell,f}$  be the set of all  $t \in T_{\ell,f}$  for which the gap associated with  $\ell$  and  $t$  is quick. Hence  $|T_\ell| \geq |T_{\ell,f}|/2 = n/(4\ell)$ . Since  $G_{\ell,t}$  cannot be larger than the number of cell probes in the gap, we have by the definition of a quick gap that for any  $t \in T_\ell$ ,

$$\mathbb{E}[G_{\ell,t}] \leq \frac{\kappa \cdot \ell \cdot \lg n}{4w}.$$

Using the above inequality together with Inequality (1) we can finally provide a non-trivial lower bound on the sum of the information transfers:

$$\begin{aligned} \sum_{\ell \in L} \sum_{t \in T_\ell} \mathbb{E}[I_{\ell,t}] &\geq \sum_{\ell \in L} \sum_{t \in T_\ell} \left( \frac{\kappa \cdot \ell \cdot \lg n}{2w} - 1 - \mathbb{E}[G_{\ell,t}] \right) \\ &\geq \sum_{\ell \in L} \sum_{t \in T_\ell} \left( \frac{\kappa \cdot \ell \cdot \lg n}{2w} - 1 - \frac{\kappa \cdot \ell \cdot \lg n}{4w} \right) \\ &\geq \frac{\kappa \cdot \lg n}{5w} \sum_{\ell \in L} \sum_{t \in T_\ell} \ell \geq \frac{\kappa \cdot \lg n}{5w} \sum_{\ell \in L} (|T_\ell| \cdot \ell) \geq \frac{\kappa \cdot \lg n}{5w} \sum_{\ell \in L} \left( \frac{n}{4\ell} \cdot \ell \right) \\ &= \frac{\kappa \cdot n \cdot \lg n}{20w} \cdot |L| \geq \frac{\kappa \cdot n \cdot \lg n}{20w} \cdot \frac{\lg n}{4 \lg \lg n} \in \Theta \left( \frac{n \cdot \lg^2 n}{w \cdot \lg \lg n} \right). \end{aligned}$$

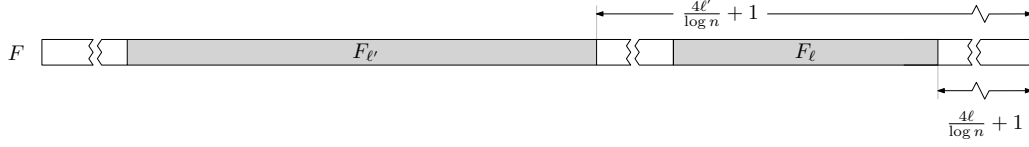


Figure 1: Two subarrays  $F_{\ell}$  and  $F_{\ell'}$  of  $F$ . There is never an overlap between any two such subarrays.

By Lemmas 1 and 2 this lower bound is also a bound on the expected total number of cell probes performed by the algorithm over the  $n$  arrivals in  $U$ . By Yao's minimax principle, as discussed in Section 2.2, this implies that any randomised algorithm on its worst case input has the same lower bound on its expected running time. The amortised time per arriving value is obtained by dividing the running time by  $n$ . This concludes the proofs of Theorem 2 and 1 .

## 4 The hard distribution for convolution

In this section we prove Lemma 4, which says that assuming Conjecture 1 holds, there exists a real constant  $\kappa > 0$  and a fixed array  $F \in \{0, 1\}^n$  such that for all  $\ell \in L$  and  $t \in [n/2]$ , when  $U$  is chosen uniformly at random from  $\{0, 1\}^n$  then the conditional entropy of the outputs  $A_{\ell,t}$  is

$$H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t}) \geq \kappa \cdot \ell \cdot \lg n,$$

for any fixed  $\tilde{u}_{\ell,t}$ . We begin by discussing the fixed array  $F$ .

### 4.1 The array $F$

For each  $\ell \in L$  there is a subarray of  $F$  of length  $\ell \lg \ell + \ell - 1$ . Each such subarray, which we denote  $F_{\ell}$ , is at distance  $4\ell / \lg n + 1$  from the right-hand end of  $F$ , which is one more than the length of the gap associated with  $\ell$ . Figure 1 illustrates two subarrays  $F_{\ell}$  and  $F_{\ell'}$ , where  $\ell, \ell' \in L$  and  $\ell < \ell'$ . By the properties discussed in Section 2.3 we know that the length of the gap associated with  $\ell'$  is larger than the length of  $F_{\ell}$  plus the length of the gap associated with  $\ell$ . Hence there is no overlap between the subarrays  $F_{\ell}$  and  $F_{\ell'}$ .

Given any of the subarrays  $F_{\ell}$  and an array  $U_{\ell}$  of length  $\ell \lg \ell$ , we write  $F_{\ell} \otimes U_{\ell}$  to denote the  $\ell$ -length array that consists of all inner products between  $U_{\ell}$  and every substring of  $F_{\ell}$ . More precisely, for  $i \in [\ell]$  the  $i$ -th component of  $F_{\ell} \otimes U_{\ell}$  is

$$(F_{\ell} \otimes U_{\ell})[i] = \sum_{j \in [\ell \lg \ell]} (F_{\ell}[i+j] \cdot U_{\ell}[j]).$$

As Figure 2 shows we may think of  $F_{\ell} \otimes U_{\ell}$  as the inner products of  $U_{\ell}$  and its aligned subarray of  $F_{\ell}$  as  $U_{\ell}$  slides along  $F_{\ell}$ . We can show that there exist subarrays  $F_{\ell}$  such

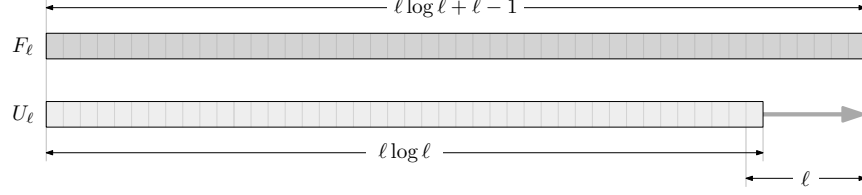


Figure 2: The  $\ell$ -length array  $F_\ell \otimes U_\ell$  contains the inner products of  $U_\ell$  and corresponding subarrays of  $F_\ell$  as  $U_\ell$  slides along  $F_\ell$ .

that the entropy of  $F_\ell \otimes U_\ell$  is high when  $U_\ell$  is drawn uniformly at random from  $\{0, 1\}^{\ell \lg \ell}$ . Observe that the entropy of  $F_\ell \otimes U_\ell$  is upper bounded by the entropy of  $U_\ell$ , which is exactly  $\ell \lg \ell$ . The proof of the next lemma is based on a conjecture that we describe shortly.

**Lemma 7.** *There exists a real constant  $\varepsilon > 0$  such that for all  $n$  and  $\ell \in L$  there is a subarray  $F_\ell$  for which the entropy of  $F_\ell \otimes U_\ell$  is at least  $\varepsilon \cdot \ell \lg \ell$  when  $U_\ell$  is drawn uniformly at random from  $\{0, 1\}^{\ell \lg \ell}$ .*

In order to finish the description of the array  $F$  we choose each subarray  $F_\ell$  such that the entropy of  $F_\ell \otimes U_\ell$  is at least  $\varepsilon \cdot \ell \lg \ell$  when  $U_\ell$  is drawn uniformly at random from  $\{0, 1\}^{\ell \lg \ell}$ , where  $\varepsilon$  is the constant of Lemma 7. Any element of  $F$  that is not part of any of the subarrays  $F_\ell$  is chosen arbitrarily. This concludes the description of the array  $F$ .

## 4.2 Proving the entropy lower bound

We are now ready to prove Lemma 4, that is lower bound the conditional entropy of  $A_{\ell,t}$ .

*Proof of Lemma 4.* Let  $F$  be the array described above and let  $U$  be drawn uniformly at random from  $\{0, 1\}^n$ . Let  $\ell \in L$  and  $t \in [n/2]$ . Thus, conditioned on any fixed  $\tilde{U}_{\ell,t}$ , the distribution of  $U_{\ell,t}$  is uniform on  $\{0, 1\}^{\ell \lg \ell}$ .

Recall that  $U_{\ell,t}$  arrives in the stream between arrival  $t_0$  and  $t_1$ , after which  $4\ell/\lg n$  values arrive in the gap. Thus, at the beginning of the second interval, at arrival  $t_2$ ,  $U_{\ell,t}$  is aligned with the  $(\ell \lg \ell)$ -length suffix of the subarray  $F_\ell$  of  $F$ . Over the  $\ell$  arrivals in the second interval,  $U_{\ell,t}$  slide along  $F_\ell$  similarly to Figure 2 (only in reversed direction of what the diagram shows). Since all values in  $\tilde{U}_{\ell,t}$  are fixed, the outputs  $A_{\ell,t}$  uniquely specify  $F_\ell \otimes U_{\ell,t}$ . Thus, by Lemma 7, the conditional entropy

$$H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t}) \geq \varepsilon \cdot \ell \cdot \lg \ell, \geq \frac{\varepsilon}{4} \cdot \ell \cdot \lg n,$$

since  $\ell \geq n^{1/4}$ . By setting the constant  $\kappa$  to  $\varepsilon/4$  we have proved Lemma 4.  $\square$

The last piece remaining is the proof of Lemma 7. This proof is based on a conjecture which we state as follows. Recall that a *Toeplitz* matrix (or upside-down *Hankel* matrix) is constant on each descending diagonal from left to right.

**Conjecture 1.** *There exist two positive real constants  $\alpha \leq 1$  and  $\gamma \leq 1$  such that for any  $h$  there is a  $(0, 1)$ -Toeplitz matrix  $M$  of height  $h$  and width  $\alpha \cdot h \lg h$  with the property that the entropy of the product  $Mv$  is at least  $\gamma \cdot h \lg h$ , where  $v$  is a column vector of length  $\alpha \cdot h \lg h$  whose elements are chosen independently and uniformly at random from  $\{0, 1\}$ .*

This conjecture might at first seem surprising as the matrix is non-square. However, an essentially equivalent statement was shown to be true for general  $(0, 1)$ -matrices in 1974 by Erdős and Spencer [7]. Moreover, even stronger statements for general  $(0, 1)$ -matrices form the basis of the well studied “Coin Weighing Problem with a Spring Scale” (see [3] and references therein). We will now show how to prove Lemma 7 using the conjecture.

*Proof of Lemma 7 (assuming Conjecture 1).* Let  $\alpha$  and  $\gamma$  be the two constants in Conjecture 1. Let  $h = \ell$  and let  $M$  be a  $(0, 1)$ -Toeplitz matrix of height  $\ell$  and width  $\alpha \cdot \ell \lg \ell$  with the property of Conjecture 1.

We now define a new matrix  $M_\ell$  of height  $\ell$  and width  $\ell \lg \ell$  such that the submatrix of  $M_\ell$  that spans the first  $\alpha \cdot \ell \lg \ell$  columns equals  $M$ . Remaining elements of  $M_\ell$  are filled in arbitrarily from the set  $\{0, 1\}$  so that  $M_\ell$  becomes Toeplitz.

Let  $v_\ell$  be a random column vector of length  $\ell \lg \ell$  such that the first  $\alpha \cdot \ell \lg \ell$  elements are chosen independently and uniformly from  $\{0, 1\}$ . The remaining elements are fixed arbitrarily. By Conjecture 1 we have that the entropy of  $M_\ell v_\ell$  is at least  $\gamma \cdot \ell \lg \ell$ . Thus, if we instead pick all elements of  $v_\ell$  independently and uniformly at random from  $\{0, 1\}$  then the conditional entropy of  $M_\ell v_\ell$ , conditioned on all but the first  $\alpha \cdot \ell \lg \ell$  elements, is at least  $\gamma \cdot \ell \lg \ell$ . Hence the entropy of  $M_\ell v_\ell$  is also at least  $\gamma \cdot \ell \lg \ell$ .

Now, since  $M_\ell$  is Toeplitz we have that the first column and the first row of  $M_\ell$  define the entire matrix. These  $\ell + \ell \lg \ell - 1$  elements can be represented with an array  $F_\ell$  of length  $\ell + \ell \lg \ell - 1$  such that  $F_\ell \otimes v_\ell = M_\ell v_\ell$ , where we abuse notation slightly. In other words, the elements of the product  $M_\ell v_\ell$  correspond exactly to the inner products obtained by sliding  $U_\ell$ , where  $U_\ell$  is the array version of the column vector  $v_\ell$ , along  $F_\ell$ . Thus, the entropy of  $F_\ell \otimes U_\ell$  is at least  $\gamma \cdot \ell \lg \ell$ . We may therefore set the constant  $\varepsilon$  in the statement of the lemma to equal the constant  $\gamma$  of Conjecture 1. This concludes the proof of Lemma 7.  $\square$

## 5 The hard distribution for $L_2$ -rearrangement

In this section we prove Lemma 5 which says that there exists a real constant  $\kappa > 0$  and a fixed array  $F \in \{0, 1\}^n$  such that for all  $\ell \in L$  and all  $t \in [n/4]$  such that  $t \bmod 4 = 0$ , when  $U$  is chosen uniformly at random from  $\{0101, 1010\}^{\frac{n}{4}}$  then the conditional entropy of the outputs  $A_{\ell,t}$  is  $H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t}) \geq \kappa \cdot \ell \cdot \lg n$  for any fixed  $\tilde{u}_{\ell,t}$ . We begin by discussing the fixed array  $F$ .



## 5.1 The array $F$

As for the convolution problem, for each  $\ell \in L$  there is a subarray of  $F$  of length  $\ell \lg \ell + \ell$ . Each such subarray, which we denote  $F_\ell$ , is at distance  $4\ell/\lg n + 1$  from the right-hand end of  $F$ . This is the same high-level structure as for the convolution problem and so again, there is no overlap between the subarrays  $F_\ell$  and  $F_{\ell'}$  and further, Figure 1 in Section 2.3 is accurate here too.

Given any of the subarrays  $F_\ell$  and an array  $U_\ell$  of length  $(\ell \lg \ell)$ , we write  $F_\ell \odot U_\ell$  to denote the  $(\ell/4)$ -length array that consists of the  $L_2$ -rearrangement distances between  $U_\ell$  and every *fourth*  $(\ell \lg \ell)$ -length substring of  $F_\ell$ . More precisely, for  $4i \in [\ell]$ , the value of  $F_\ell \odot U_\ell[i]$  is the  $L_2$ -rearrangement distance between  $F_\ell[4i, 4i + \ell \lg \ell - 1]$  and  $U_\ell$ .

Our main focus in this section is on proving Lemma 8 which can be seen as an analogue of Lemma 5 for a fixed length of  $\ell$ :

**Lemma 8.** *There exists a real constant  $\varepsilon > 0$  such that for all  $n$  and  $\ell \in L$  there is a subarray  $F_\ell$  for which the entropy of  $F_\ell \odot U_\ell$  is at least  $\varepsilon \cdot \ell \lg \ell$  when  $U_\ell$  is drawn uniformly at random from  $\{0101, 1010\}^{\frac{\ell}{4} \lg \ell}$ .  $F_\ell$  contains an equal number of 0s and 1s.*

In order to finish the description of the array  $F$  we choose each subarray  $F_\ell$  in accordance with Lemma 8. Any region of  $F$  that is not part of any of the subarrays  $F_\ell$  is filled with repeats of ‘01’. This ensures that these regions contain an equal number of zeros and ones (it is easily verified that each region has an even length). This concludes the description of the array  $F$ .

The proof of Lemma 5 then follows from Lemma 8. It is conceptually similar to the proof of Lemma 4 for convolution which follows from Lemma 7. However, for the convolution problem, our proof relied on the fact that we could essentially consider the convolution of each pair  $F_\ell$  and  $U_\ell$  separately and simply add them up to find the convolution of  $F$  and  $U$ . This is less immediate for  $L_2$ -rearrangement distance because we need to rule out the possibility of characters from some  $U_\ell$  being moved to positions in  $F_{\ell'}$  for  $\ell \neq \ell'$ . The proof (and the lower bound in general) relies on a key property of  $L_2$ -arrangement (proven in Lemma 3.1 from [1]) which states that under the optimal rearrangement permutation, the  $i$ -th one (resp. zero) in one string is moved to the  $i$ -th one (resp. zero) in the other. By controlling how the zeroes and ones are distributed in  $U$  and  $F$ , we can limit how far any character is moved.

We are now ready to prove Lemma 5, that is lower bound the conditional entropy of  $A_{\ell,t}$ .

*Proof of Lemma 5.* Let  $F$  be the array described above and let  $U$  be drawn uniformly at random from  $\{0101, 1010\}^{\frac{n}{4}}$ . Let  $\ell \in L$  and  $t \in [n/2]$ . Thus, conditioned on any fixed  $\tilde{U}_{\ell,t}$ , the distribution of  $U_{\ell,t}$  is uniform on  $\{0101, 1010\}^{\frac{\ell}{4} \lg \ell}$ .

Recall that  $U_{\ell,t}$  arrives in the stream between arrival  $t_0$  and  $t_1$ , after which  $4\ell/\lg n$  values arrive in the gap. Thus, at the beginning of the second interval, at arrival  $t_2$ ,  $U_{\ell,t}$

is aligned with the  $(\ell \lg \ell)$ -length suffix of the subarray  $F_\ell$  of  $F$ . Over the  $\ell$  arrivals in the second interval,  $U_{\ell,t}$  slides along  $F_\ell$  similarly to Figure 2. We now prove that since all values in  $\tilde{U}_{\ell,t}$  are fixed, the outputs  $A_{\ell,t}$  uniquely specify  $F_\ell \odot U_{\ell,t}$ . The analogous property for convolution was immediate. First observe that by construction the prefix of  $F$  up to the start of  $F_\ell$  contains an equal number of 0s and 1s. Similarly for  $F_\ell$  itself and the suffix from  $F_\ell$  to the end of  $F$ . Once in every four arrivals, the substring of  $U$  aligned with  $F$  is guaranteed (by construction) to also have an equal number of 0s and 1s. Therefore the  $L_2$  rearrangement distance is finite. It was proven in Lemma 3.1 from [1] that (rephrased in our notation) under the optimal rearrangement permutation, the  $k$ -th one (resp. zero) in  $F$  is moved to the  $k$ -th one (resp. zero) in  $U$ . Therefore, every element of  $U_\ell$  is moved to an element in  $F_\ell$ . We can therefore recover any output in  $F_\ell \odot U_{\ell,t}$  by taking the corresponding output in  $A_{\ell,t}$  and subtracting, the costs of moving the elements that are in  $U$  but not in  $U_\ell$ . It is easily verified that as  $t$  is divisible by four, the corresponding output in  $A_{\ell,t}$  is one of those guaranteed to have an equal number of 0s and 1s. Thus, by Lemma 8, the conditional entropy

$$H(A_{\ell,t} \mid \tilde{U}_{\ell,t} = \tilde{u}_{\ell,t}) \geq \varepsilon \cdot \ell \cdot \lg \ell, \geq \frac{\varepsilon}{4} \cdot \ell \cdot \lg n,$$

since  $\ell \geq n^{1/4}$ . By setting the constant  $\kappa$  to  $\varepsilon/4$  we have proved Lemma 5.  $\square$

## 5.2 The proof of Lemma 8

In this section we prove Lemma 8. We begin by explaining the high-level approach which will make one final composition of both  $F_\ell$  and  $U_\ell$  into sub-arrays. For any  $j \geq 0$ , let  $U_\ell^j = U_\ell[\ell \cdot j, \ell \cdot (j+1) - 1]$  i.e.  $U_\ell^j$  is the  $j$ -th consecutive  $\ell$ -length sub-array of  $U_\ell$ . The key property that we will prove in this section is given in Lemma 9 which intuitively states that given half of the bits in  $U_\ell$ , we can compute the other half with certainty.

**Lemma 9.** *Let  $U_\ell$  be chosen arbitrarily from  $\{0101, 1010\}^{\ell/4}$ . Given  $F_\ell$ ,  $F_\ell \odot U_\ell$  and  $U_\ell^{2j+1}$  for all  $j \geq 0$ , it is possible to uniquely determine  $U_\ell^{2j}$  for all  $j \geq 0$ .*

Before we prove Lemma 9, we briefly justify why Lemma 8 is in-fact a straight-forward corollary of Lemma 9. If we pick  $U_\ell$  uniformly at random from  $\{0101, 1010\}^{\ell/4}$  then by Lemma 9, the conditional entropy,  $H(F_\ell \odot U_\ell \mid U_\ell^{2j+1} \text{ for all } j)$  is  $\Omega(\ell \lg \ell)$ . This is because we always recover  $\Theta(\lg \ell)$  distinct  $U_\ell^{2j}$ , each of which is independent and has entropy  $\Omega(\ell)$  bits. It then immediately follows that  $H(F_\ell \odot U_\ell) \geq H(F_\ell \odot U_\ell \mid U_\ell^{2j+1} \text{ for all } j)$  as required. We also require for Lemma 8 that  $F_\ell$  contains an equal number ones and zeros. This follows immediately from the description of  $F_\ell$  below.

**The subarray  $F_\ell$**  We now give the description of  $F_\ell$  which requires one final decomposition into subarrays. This description is supported by Figure 3. For each  $j \in \llbracket (\lg \ell)/2 \rrbracket$ ,

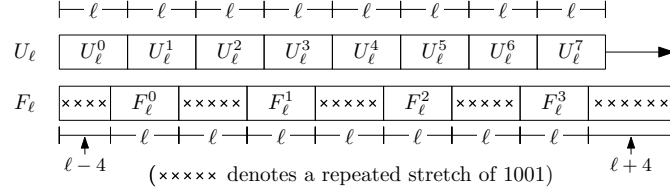


Figure 3: The high-level structure of  $U_\ell$  and  $F_\ell$ .

$F_\ell$  contains a subarray  $F_\ell^j$  of length  $\ell$ . Intuitively, each sub-array  $F_\ell^j$  will be responsible for recovering  $U_\ell^{2j}$ . These subarrays occur in order in  $F_\ell$ . Before and after each  $F_\ell^j$  there are stretches of repeats of the string 1001. Specifically, before  $F_\ell^1$  there are  $\ell/4 - 1$  repeats of the string 1001. Between each  $F_\ell^j$  and  $F_\ell^{j+1}$  there are  $\ell/4$  repeats of the string 1001 and after  $F_\ell^{\lfloor (\lg \ell)/2 \rfloor - 1}$  there are  $\ell/4 + 1$  repeats. These repeats of 1001 are simply for structural padding and as we will see the contribution of these repeated 1001 strings to the  $L_2$  rearrangement distance is independent of  $U_\ell$ . This follows because the cost of rearranging 1001 to match either 1010 or 0101 is always 2.

The structure of  $F_\ell^j$  is as follows  $F_\ell^j = 10^{(2^j+3)}1^{(\ell/4-1)}0^{(\ell/4-(2^j+3))}$ . Here  $0^z$  (resp.  $1^z$ ) is a string of  $z$  zeros (resp. ones). Intuitively, the reason that the stretch of 0s at the start of  $F_\ell^j$  is exponentially increasing with  $j$  is so that distance the second one in  $F_\ell^j$  (immediately after the stretch of 0s) is forced to move also exponentially increasing with  $j$ . This will allow us to recover each  $U_\ell^{2j}$  from a different bit in the outputs.

**Proof of Lemma 9** We are now in a position to prove Lemma 9. Our main focus will be on first proving that given  $F_\ell$ ,  $U_\ell^{2j+1}$  for all  $j$  and  $F_\ell \odot U_\ell$ , we can uniquely determine  $U_\ell^{2j}[\ell - 4, \ell - 1]$  for each  $j \geq 0$ . That is, for each  $j$  whether the last four symbols of  $U_\ell^{2j}$  are 0101 or 1010. This is shown diagrammatically in Figure 4. We argue that by a straight-forward repeated application of this argument we can in-fact recover the whole of  $U_\ell^{2j}$  for all  $j \geq 0$ .

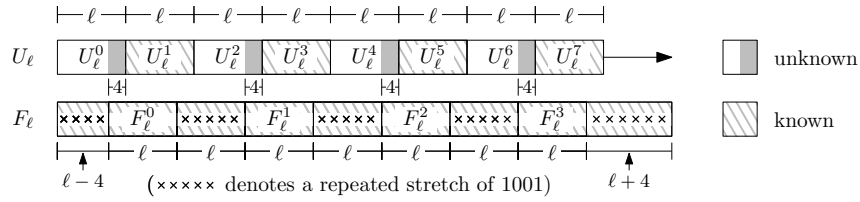


Figure 4: We can determine  $U_\ell^{2j}[\ell - 4, \ell - 1]$  if we know  $F_\ell$ , every  $U_\ell^{2j+1}$  and  $F_\ell \odot U_\ell$ .

We will begin by making some simplifying observations about  $(F_\ell \odot U_\ell)[0]$ . Recall that  $(F_\ell \odot U_\ell)[0]$  was defined to be the  $L_2$  rearrangement distance between  $F_\ell[0, |U_\ell| - 1]$  and

$U_\ell$ . The first observation is that is finite because both strings contain an equal number of zeros and ones.

The  $L_2$  rearrangement distance  $(F_\ell \odot U_\ell)[0]$  can be expressed as the sum of the *contributions* from moving each  $U_\ell[i]$ , over all  $i \in [m]$ . Let the contribution of  $U_\ell[i]$ , denoted,  $\text{CT}(i)$  be the square of the distance that  $U_\ell[i]$  is moved under the optimal rearrangement. We then have that  $(F_\ell \odot U_\ell)[0] = \sum_i \text{CT}(j)$ . Finally, we let  $D^*$  be the sum of the contributions of the locations in every  $U_\ell^{2j}[\ell - 4, \ell - 1]$ , i.e.  $D^* = \sum_j \sum_{k=0}^3 (\text{CT}(2j \cdot \ell + (\ell - 4) + k))$ . We will also refer to the contribution of a substring which is defined naturally to be the sum of the contributions of its constituent characters. For example the contribution of  $U_\ell^j$  equals  $\sum \{\text{CT}(r) \mid r \in [\ell \cdot j, \ell \cdot (j + 1) - 1]\}$ .

Our proof will be in two stages. First we will prove in Lemma 10 that we can compute  $D^*$  from  $F_\ell$ ,  $F_\ell \odot U_\ell$  and  $U_\ell^{2j+1}$  for all  $j \geq 0$ . Second we will prove that for any  $j > 0$ , we can determine  $U_\ell^{2j}[\ell - 4, \ell - 1]$  from  $D^*$ .

In the proof of Lemma 10 we argue that  $D^*$  can be calculated directly from  $(F_\ell \odot U_\ell)[0]$  by subtracting the contributions of  $U_\ell^{2j+1}$  and  $U_\ell^{2j}[0, \ell - 5]$  for all  $j \geq 0$ . More specifically, we will prove that the contribution of any  $U_\ell^{2j+1}$  can be calculated from  $U_\ell^{2j+1}$  and  $F_\ell$ , which are both known. In particular, the contribution of any  $U_\ell^{2j+1}$  is independent of every unknown  $U_\ell^{2j}$ . Further, we will prove that although  $U_\ell^{2j}$  is unknown, the contribution of  $U_\ell^{2j}[0, \ell - 5]$ , always equals  $\ell/2 - 2$ , regardless of the choice of  $U_\ell$ .

**Lemma 10.**  $D^*$  can be computed from  $F_\ell$ ,  $F_\ell \odot U_\ell$  and  $U_\ell^{2j+1}$  for all  $j \geq 0$ .

*Proof.* The value of  $D^*$  is calculated directly from  $(F_\ell \odot U_\ell)[0]$  by subtracting the contributions of  $U_\ell^{2j+1}$  and  $U_\ell^{2j}[0, \ell - 5]$  for all  $j \geq 0$ . We will now prove that for any  $j$ , we can calculate the contribution of  $U_\ell^{2j+1}$  and that contribution of  $U_\ell^{2j}[0, \ell - 5]$ , in-fact always equals  $\ell/2 - 2$ , regardless of the choice of  $U_\ell$ .

In this proof we rely heavily on Lemma 3.1 from [1] which states that under the optimal rearrangement permutation, the  $i$ -th one (resp. zero) in  $U_\ell$  is moved to the  $i$ -th one (resp. zero) in  $F_\ell[0, |U_\ell| - 1]$ .

For any  $j$ , consider,  $U_\ell^{2j}$  and  $U_\ell^{2j+1}$ . The number of ones in  $U_\ell^{2j}$  (resp.  $U_\ell^{2j+1}$ ) is fixed, independent of the choice of  $U_\ell$ . In particular there are exactly  $\ell/2$  zeroes and  $\ell/2$  ones. It is easily verified that, by construction,  $F_\ell[2j \cdot \ell, (2j + 2) \cdot \ell - 1]$  also contains exactly  $\ell$  zeros and  $\ell$  ones. Therefore, the  $i$ -th one (resp. zero) in  $U_\ell^{2j}$  is moved to the  $i$ -th one (resp. zero) in  $F_\ell[2j \cdot \ell, (2j + 2) \cdot \ell - 1]$ . Similarly, the  $i$ -th one (resp. zero) in  $U_\ell^{2j}$  is moved to the  $(i + \ell/2)$ -th one (resp. zero) in  $F_\ell[2j \cdot \ell, (2j + 2) \cdot \ell - 1]$ .

Consider any  $U_\ell^{2j+1}$  which is known. By the above observation, we can therefore determine which position in  $F_\ell[2j \cdot \ell, (2j + 2) \cdot \ell - 1]$ , each character in  $U_\ell^{2j+1}$  is moved to under the optimal rearrangement. From this we can then directly compute the contribution of each  $U_\ell^{2j+1}$  to  $(F_\ell \odot U_\ell)[0]$ .

Consider any  $U_\ell^{2j}$  which is unknown. As observed above, the  $i$ -th one (resp. zero) in  $U_\ell^{2j}$  is moved to the  $i$ -th one (resp. zero) in  $F_\ell[2j \cdot \ell, (2j + 2) \cdot \ell - 1]$ . By construction, we

have that  $F_\ell[2j \cdot \ell, (2j+1) \cdot \ell - 5]$  consists entirely of repeats of 1001. Further for any  $i$ , we have that  $U_\ell^{2j}[4i, 4i+3]$  is either 1010 or 0101. Therefore for all  $i < \ell/4$  we have that the two ones (resp. zeroes) in  $U_\ell^{2j}[4i, 4i+3]$  are moved to the two ones (resp. zeroes) in  $F_\ell[2j \cdot \ell + 4i, 2j \cdot \ell + 4i + 3] = 1001$ . The key observation is that regardless of whether  $U_\ell^{2j}[4i, 4i+3] = 1010$  or 0101, the contribution of  $U_\ell^{2j}[4i, 4i+3]$  is 2. Therefore for any  $U_\ell$ , the contribution of  $U_\ell^{2j}[0, \ell - 5]$  is  $\ell/2 - 2$ .  $\square$

In Lemma 11 we will prove that we can compute  $U_\ell^{2j}[\ell - 4, \ell - 1]$  from  $D^\star$  (for any sufficiently large  $j$ ). The intuition behind this is given by Fact 1 which gives an explicit formula for the contribution of  $U_\ell^{2j}[\ell - 4, \ell - 1]$ . Observe that the contribution depends only on whether  $U_\ell^{2j}[\ell - 4, \ell - 1]$  equals 1010 ( $v_j = 1$ ) or 0101 ( $v_j = 0$ ). The intuition is that we can extract  $v_j$  from the  $(j+1)$ -th bit of  $D^\star$ .

**Fact 1.** *For any  $j$ , let  $v_j = 1$  if  $U_\ell^{2j}[\ell - 4, \ell - 1] = 1010$  and  $v_j = 0$  otherwise. The contribution of  $U_\ell^{2j}[\ell - 4, \ell - 1]$  is exactly  $v_j \cdot 2^{j+1} + 2^{2j} + 2$ .*

*Proof.* We begin by arguing that under the optimal rearrangement permutation, the two ones (resp. zeroes) in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  are moved to the leftmost two ones (resp. zeroes) in  $F_\ell^j$ . We will again rely heavily on Lemma 3.1 from [1] which states that under the optimal rearrangement permutation, the  $i$ -th one (resp. zero) in  $U_\ell$  is moved to the  $i$ -th one (resp. zero) in  $F_\ell[0, |U_\ell| - 1]$ .

In the proof of Lemma 10, we argued that the  $i$ -th one (resp. zero) in  $U_\ell^{2j}$  is moved to the  $i$ -th one (resp. zero) in  $F_\ell[2j \cdot \ell, (2j+2) \cdot \ell - 1]$ . It is easily verified that  $F_\ell[2j \cdot \ell, (2j+2) \cdot \ell - 1]$  consists of exactly  $\ell/4 - 1$  repeats of 1001 followed by  $F_\ell^j$ . Therefore as  $U_\ell^{2j}[0, \ell - 5]$  contains exactly  $\ell/2 - 2$  ones (resp. zeroes), the two ones in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  are indeed moved to the leftmost two ones (resp. zeroes) in  $F_\ell^j$ .

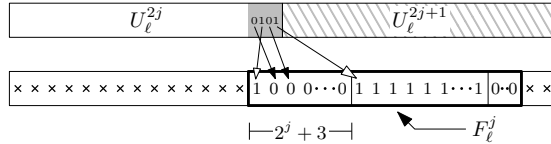


Figure 5: The rearrangement of the symbols in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  under the optimal rearrangement permutation. The highlighted region is  $F_\ell^j$ .

We now argue about the contribution of each character in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  in turn. This argument is supported by Figure 5.

First consider the first one in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  which is moved to the first one in  $F_\ell^j$ . By construction we have that, when indexed from the start of  $F_\ell$ , the leftmost one in  $F_\ell^j$  is at position  $x = (2j+1) \cdot \ell - 4$ . Similarly, when indexed from the start of  $U_\ell$ , the leftmost one

in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  is at position  $x + (1 - v_j)$ . Therefore the contribution of the first one in  $U_\ell^{2j}$  is exactly  $(1 - v_j)^2$ .

The first zero in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  is moved to the first zero in  $F_\ell^j$ . The first zero in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  is at position  $x + v_j$ . The first zero in  $F_\ell^j$  is at position  $x + 1$ . Therefore the contribution of the first zero in  $U_\ell^{2j}$  is also exactly  $(1 - v_j)^2$ .

The second zero in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  is moved to the second zero in  $F_\ell^j$ . The second zero in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  is at position  $x + 2 + v_j$  and the second zero in  $F_\ell^j$  is at position  $x + 2$ . Therefore the contribution of the first zero in  $U_\ell^{2j}$  is exactly  $v_j^2$ .

Finally, the second one in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  is moved to the second one in  $F_\ell^j$ . Again, by construction, we have that the second one in  $F_\ell^j$  is at position  $x + 3 + 2^j$  (indexed from the start of  $F_\ell$ ). Similarly, the second one in  $U_\ell^{2j}[\ell - 4, \ell - 1]$  is at position  $x + 3 - v_j$ . Therefore the contribution of the second one in  $U_\ell^{2j}$  is exactly,  $(2^j + v_j)^2$ . Summing over all four characters we have that the total contribution of  $U_\ell^{2j}[\ell - 4, \ell - 1]$  is:

$$(2^j + v_j)^2 + v_j^2 + 2(1 - v_j)^2$$

Expanding and simplifying we have that this is

$$(2^{2j} + 2) + v_j \cdot 2^{j+1} + 4v_j^2 - 4v_j$$

As  $v_j \in \{0, 1\}$  we have that  $v_j^2 = v_j$  so this simplifies further to:

$$v_j \cdot 2^{j+1} + 2^{2j} + 2$$

□

We can now prove Lemma 11 which follows almost immediately from Fact 1.

**Lemma 11.** *For any  $j \geq 0$ , it is possible to compute  $U_\ell^{2j}[\ell - 4, \ell - 1]$  from  $D^*$ .*

*Proof.* Let  $D_2^*$  equal  $D^* - \sum_j (2^{2j} + 2)$  which can be calculated directly from  $D^*$ . An alternative and equivalent definition of  $D_2^*$  follows from Fact 1 and is given by

$$D_2^* = \sum_j v_j \cdot 2^{j+1}.$$

We can therefore compute  $v_j$  and hence  $U_\ell^{2j}[\ell - 4, \ell - 1]$  by inspecting the  $(j + 1)$ -th bit of  $D_2^*$ .

□

**Recovering the rest of  $U_{\ell,(2j)}$**  So far we have only proven that we can recover  $U_{\ell}^{2j}[\ell - 4, \ell - 1]$  for all  $j$ . The claim that we can in-fact recover the whole of  $U_{\ell}^{2j}$  follows by repeatedly application of the the proof above. Specifically, once we have recovered  $U_{\ell}^{2j}[\ell - 4, \ell - 1]$  for all  $j$ , we can use this additional information (and  $(F_{\ell} \odot U_{\ell})[1]$  instead of  $(F_{\ell} \odot U_{\ell})[0]$ ) to recover  $U_{\ell}^{2j}[\ell - 8, \ell - 5]$  for all  $j$  and so on. More formally we proceed by induction on increasing  $k$  by observing that using the above argument given  $F_{\ell}$ ,  $(F_{\ell} \odot U_{\ell})[k]$ ,  $U_{\ell}^{2j+1}$  for all  $j \geq 0$  and  $U_{\ell}^{2j+1}[\ell - 4k, \ell - 1]$  for all  $j \geq 0$  we can recover  $U_{\ell}^{2j+1}[\ell - 4k - 4, \ell - 4k - 1]$  for all  $j$ .

## References

- [1] A. Amir, Y. Aumann, G. Benson, A. Levy, O. Lipsky, E. Porat, S. Skiena, and U. Vishne. Pattern matching with address errors: rearrangement distances. In *SODA '06: Proc. 17<sup>th</sup> ACM-SIAM Symp. on Discrete Algorithms*, pages 1221–1229. ACM Press, 2006.
- [2] A. Amir, Y. Aumann, G. Benson, A. Levy, O. Lipsky, E. Porat, S. Skiena, and U. Vishne. Pattern matching with address errors: Rearrangement distances. *Journal of Computer System Sciences*, 75(6):359–370, 2009.
- [3] N. H. Bshouty. Optimal algorithms for the coin weighing problem with a spring scale. In *COLT '09: Proc. 22<sup>nd</sup> Annual Conference on Learning Theory*, 2009.
- [4] R. Clifford, M. Jalsenius, and B. Sach. Cell-probe bounds for online edit distance and other pattern matching problems. In *SODA '15: Proc. 26<sup>th</sup> ACM-SIAM Symp. on Discrete Algorithms*, 2015.
- [5] R. Clifford and M. Jalsenius. Lower bounds for online integer multiplication and convolution in the cell-probe model. In *ICALP '11: Proc. 38<sup>th</sup> International Colloquium on Automata, Languages and Programming*, pages 593–604, 2011.
- [6] R. Clifford, M. Jalsenius, and B. Sach. Tight cell-probe bounds for online hamming distance computation. In *SODA '13: Proc. 24<sup>th</sup> ACM-SIAM Symp. on Discrete Algorithms*, pages 664–674, 2013.
- [7] P. Erdős and J. Spencer. *Probabilistic Methods in Combinatorics*. Academic Press Inc, 1974.
- [8] M. Fredman. Observations on the complexity of generating quasi-Gray codes. *SIAM Journal on Computing*, 7(2):134–146, 1978.
- [9] M. Fredman and M. Saks. The cell probe complexity of dynamic data structures. In *STOC '89: Proc. 21<sup>st</sup> Annual ACM Symp. Theory of Computing*, pages 345–354, 1989.

- [10] K. G. Larsen. The cell probe complexity of dynamic range counting. In *STOC '12: Proc. 44<sup>th</sup> Annual ACM Symp. Theory of Computing*, pages 85–94, 2012.
- [11] K. G. Larsen. Higher cell probe lower bounds for evaluating polynomials. In *FOCS '12: Proc. 53<sup>rd</sup> Annual Symp. Foundations of Computer Science*, pages 293–301, 2012.
- [12] M. Minsky and S. Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, 1969.
- [13] M. Pătraşcu. *Lower bound techniques for data structures*. PhD thesis, MIT, 2008.
- [14] M. Pătraşcu and E. D. Demaine. Tight bounds for the partial-sums problem. In *SODA '04: Proc. 15<sup>th</sup> ACM-SIAM Symp. on Discrete Algorithms*, pages 20–29, 2004.
- [15] M. Pătraşcu and E. D. Demaine. Logarithmic lower bounds in the cell-probe model. *SIAM Journal on Computing*, 35(4):932–963, 2006.
- [16] A. C.-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *FOCS '77: Proc. 18<sup>th</sup> Annual Symp. Foundations of Computer Science*, pages 222–227, 1977.
- [17] A. C.-C. Yao. Should tables be sorted? *Journal of the ACM*, 28(3):615–628, 1981.